W. Johnston
Lawrence Berkeley National Lab and NASA Ames
J. Brooke
Manchester Computing, University of Manchester

Core Functions for Production Grids

(alt. title: Minimal Grid Functions for Production Grids)

Working Draft, Version 2

Status of this Memo: This memo is a draft for technical recommendations. Distribution is unlimited.

Abstract

This document defines ***current practice*** for set of Grid services that are intended to be the ***minimal*** set of Grid functions needed for a functioning, production Grid. We call these the Core Grid Functions.

This version of the document is a draft, and is a strawman to raise the issues that will generate the discussion needed to agree on the Core Grid Functions.

Table of Contents

1. Introduction 2
2. Criteria for Minimality and Current Practice 3
3. Core Grid Functions 5
3.1 Resource Discovery 5
3.2 Resource Scheduling 7
3.3 Uniform Computing Access 9
3.4 Uniform Data Access 11
3.4.1 Flat File / Unstructured Object Access 11
3.4.2 Relational Data Base access 12
3.4.3 Object Oriented Data Base access 13
3.5 Asynchronous Information Sources (Events, Monitoring, Logging, etc.) 14
3.6 Remote Authentication, Delegation, and Secure Communication 16
3.6.1 Certification Authority and Certificate Management 17
3.6.2 User Key Management 18
3.6.3 Mutual Authentication 18
3.6.4 Secure Communication 19
3.6.5 Delegation 20
3.6.6 GSS-API 20
3.6.7 The Overall Grid Security Infrastructure Service 21
3.7 System Management and Access 23
3.8 Architectural Constraints 24
3.9 Bindings 25
3.10 Other / Future Services as Core Grid Functions 26
3.10.1 Abstraction of Computing Resource Architecture 26

3.10.2 Transactional Messaging 26
3.10.3 Reliable, Secure Multicast 26
3.10.4 Checkpoint / Restart / Coordinated Recovery 26
3.10.5 Structured Data Access 26
3.10.6 Quality of Service 26
3.10.7 Debug 26
3.10.8 Communications channel “tapping” 27
3.10.9 Authorization 27
4. Security Considerations 28
5. Glossary 28
6. Author Contact Information 28
7. Acknowledgements 28
8. Notices 28
8.1 Intellectual Property Statement 28
8.2 Full Copyright Notice 29
9. Notes and References 29

1. Introduction

This document is intended to define the ***current practice*** for a ***minimal*** set of Grid functions that provide uniform interfaces to architecturally, geographically, and administratively heterogeneous computing, data, and instrument systems that are managed by production Grids. By “production Grids” we mean the Grids that are trying to use Grid technology to provide services to a diverse user community to whom the operators of the Grid are responsible for providing a reliable and useful service. Defining these minimal services is very important because they represent the fundamental persistent infrastructure of the Grid. We use the term Core Grid Functions to represent this collection of Grid services that provide the persistent and most basic functionality of Grids.

This minimal set of functions are the services that provide the resource independence that will make the Grid a common infrastructure for all higher-level services. That is, they are the smallest set of services that are needed to build all other Grid frameworks, middleware, and applications. The minimal services may vary somewhat depending on the type of Grid resource – computing, data, instrument, etc.

Defining a “minimal” set of functions is important because:

1) This set provides a metric about whether a system is a Grid enabled system, or not. Without the minimal set of functions, there will be Grid middleware, frameworks, and applications that cannot not function correctly, or at all, on such a system without adding the missing services.
2) Since the minimal set has to be installed and managed on every system that is a Grid resource they represent most of the operational effort in building and managing Grids.

Our current understanding of Minimal Grid Functions includes

- resource discovery
- resource scheduling
- uniform computing access
- uniform data access
- asynchronous information sources (events and monitoring)
- authentication, delegation, and secure communication
- identity certificate management

- system management and access

All of these functions have elements that must be installed and managed on the Grid resources / systems, or have independent servers that provide the functions (e.g. discovery and identity certificates) and some may require both (e.g. asynchronous information sources require both services for generating events and a separate registry service).

For each of the minimal set of functions, then we identify what are the minimal characteristics that must be standardized:

- functionality
- data structures (minimal?)
- bindings (minimal?)

The relationship of the Core Grid Functions to other elements of the Grid is indicated in Figure 1.

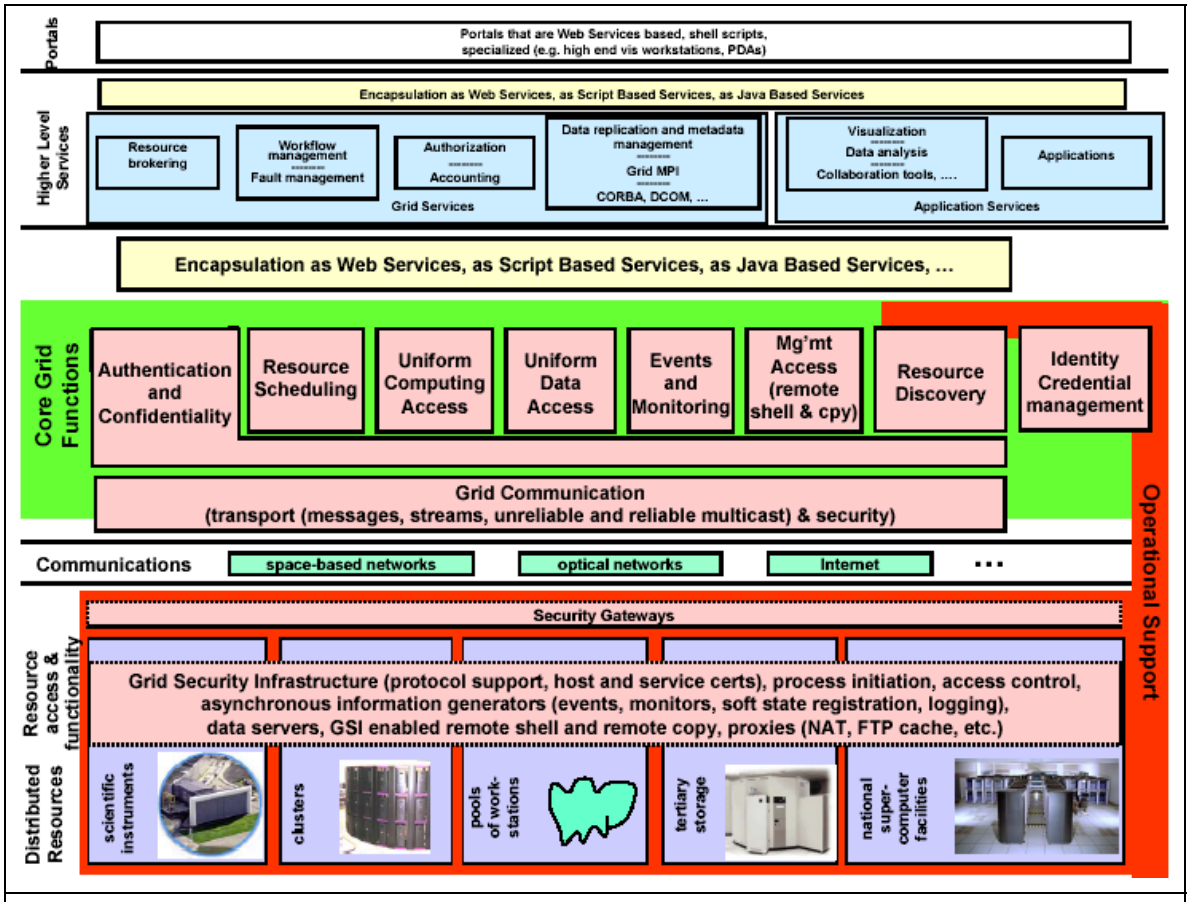

Figure 1. Core Grid Functions Relationships

2. Criteria for Minimality and Current Practice

In general, the criteria for a Core Grid Function is that it

- cannot be built on top of other Grid services,
- is essential for building other Grid services and applications, or for providing scalability or security,
- must be self contained (except possibly with respect to security).

Grid data and instrument resources may not have all of the functionality of computing resource, and may, therefore, not support a full set of CGF (e.g., Grid based scientific instruments and tertiary storage systems may (probably will not) allow job initiation).

These criteria are evaluated for the candidate list of GCS below.

This document is also intended to address “current practice.” That is, there are groups that are building, or attempting to build, Grids and use these Grids to deliver services to well defined (though possibly diverse) user community. Most of these services exist today in one form or another. However, in a few cases the implementation work is turning up a few Grid functions that are missing, and functions that we understand well enough that we believe that they can be defined and specified relatively easily. The functions will be included here as well, since we believe that they are part of the minimal set.

Issue: Are there certain bindings that should be required on the client side as part of the minimal function definition?

3. Core Grid Functions

3.1 Resource Discovery

A Grid information service must provide information about all Grid resources, and should minimize the number of persistent information servers that have to be managed in order to enable Grid services and applications. In other words, it provides the basic persistent state mechanism for the Grid.

Functionality

The discovery service should

- Provide for locating all Grid resources with specified properties, within a certain scoping, and without relying on user maintained enumerated lists
- Accommodate a dynamic resource base
- Be extensible to "all" Grid persistent state – that is, all Grid services can be sources of information, and if this information needs to be referenced and/or discovered, it should be possible to store and/or represent it in the Grid information service. E.g.
  - Virtual Organization membership
  - Computing resources
    - Available software
    - Current user allocation
  - Asynchronous Information Sources registry and data content
- Accommodate data from users, Virtual Organizations, applications

Essential characteristics

- query
- query access control
- soft state data entry
- data entry access control
- search scoping
  - hierarchical
  - "views" (e.g. MDS-2 index server)`
- aggregated queries
  - query all of the information sources within the scope of a information service

Issues

- Are aggregated queries – query all of the information sources within the scope of a information service – essential for scalability?

QoS

- query response time
- data entry limitations

Support Required on Grid Resource Platforms

- Soft state registration mechanism

Environmental Support Required

- Typically one or more servers must be managed at a site and/or in a VO

Is this a minimal service?

Yes, minimal:

- Discovery is an essential Grid function. Without discovery, you cannot build virtual systems from dynamically changing pools of resources.
- Management of persistent servers is operationally expensive, therefore for VOs and Grids to be operationally successful, it is critical to minimize the number of servers needed for a persistent Grid. Storing / representing all manner of persistent Grid information with one service is important to minimize required operational support.

No, not minimal:

Cannot address at this time:

Current Experience

- Globus MDS-2 [1]
- Current GGF docs

3.2 Resource Scheduling

Scheduling is separate from process initiation. It may involve different types of resources, some of which do not involve processes.

Functionality

- Establish a given, on-demand, virtual system relationship among an administratively independent set of Grid resources – that is, among all of the components that need to interoperate (e.g. parallel, pipelined, multi-level) to accomplish a task in the Grid environment that involves many coordinated elements
- Return information sufficient for negotiation of a common QoS (e.g. time slot) among independent resources

Characteristics

- Co-scheduling by negotiation
- Time-of-day reservation
  - Hard (e.g. at a fixed time)
  - Soft (e.g. at some time within a specified interval)
- Accounting
- Security
  - Access control based on the Grid DN (distinguished name – see section 3.6. "Remote Authentication, Delegation, and Secure Communication)

Issues

- How is a reservation guaranteed
- How to control who gets to reserve a resource

QoS

- This is a key function for implementing QoS

Support Required on Grid Resource Platforms

A scheduler operating on the resource must

- Provide time of day reservation
- Evaluate the future availability of a reservation request and pass that information back to the requester
- Support soft reservations to allow time for an external broker to negotiate a common reservation among several resources

Environmental Support Required

- none

Is this a minimal service?

Yes, minimal:

- Essential for QoS
- Not possible to emulate

No, not minimal:

Cannot address at this time:

Current Experience

- GGF
  - SchedWD8 - "Super Scheduler Steps/Framework", J. Schopf, 7/01 - overview of current user practices for scheduling across administrative domains. [2]
- Globus
  - SNAP/GARA –" SNAP: A Protocol for Negotiating Service Level Agreements and Coordinating Resource Management in Distributed Systems" - [3]
- Resource schedulers
  - PBSPro [4]
  - Maui Silver [5]

3.3 Uniform Computing Access

Most current Grid access to computing resources involves sending a script to the remote systems and then either exec-ing a shell and passing the script to the shell, or, more commonly for a purely computing resource, submitting the script to a batch queuing system.

[[John Brooke is working on this section. – wej]]

Functionality

- Initiate a process or task script on a remote Grid system
- Support queries about queue types
- Support submission to names queues (different classes of service)
- Perform access control based on Grid DN

Characteristics

Issues

- There are both practical and actual issues that arise with this service that may make different Grid computing resources appear to have a uniform interface, but they will actually be heterogeneous in that a task sequence that works on one system will not work on another due to the system configuration/architecture.
  - Most of these configuration issues relate to either setting up the environment needed to execute the Grid job, or in pre- and post staging data.
  - The issues with staging data derive from the fact that cluster-like systems (which is to say, most of today's large computing resources) may
    + not provide all services on every node. For example,
      - some nodes may not have host certificates, and therefore services requiring host certs (like GridFTP server) will not work.
      - some nodes may not have access to the Internet
  - Have different approaches to how the file system environment for a job/task is established
  - Have different ways to stage and store files locally that must be accessed by computing nodes of the resource
- Some task models require pre-staging data independently of the management of the computation – how do we deal with this in general?

QoS

- Provided by the scheduling service

Support Required on Grid Resource Platforms

- Transfer the Grid DN of users that initiate jobs into the accounting records for those jobs
- Provide an access control mechanism based on the Grid DN
- Grid process initiation service must be installed and managed
- Access control mechanisms must be maintained

Environmental Support Required

- none

Is this a minimal service?

Yes, minimal:

An essential Grid service.

No, not minimal:

Cannot address at this time:

Current Experience

- Globus gatekeeper, GRAM [6]
- Current GGF docs

3.4 Uniform Data Access

Today the primary Grid data access service is access to named, unstructured objects (e.g. "flat" files). That is, objects / files whose structure is understood only by the application that reads the files, and not by the storage system. Hence, the primary current model for Grid data access is FTP.

Two other situations are being used and/or considered in Grid storage resources:

- Support for some mechanism of sub-setting or filtering data before it leaves the storage resource.
- providing access to relational databases.

Issue: How many core access types are there?
Current experience is with objects / files, relational DBMS, and (maybe) Object Oriented DBMS.

3.4.1 Flat File / Unstructured Object Access

Functionality

- Storage access abstraction
- Partial file access
- Integrated Grid security infrastructure security and access control based on the Grid DN

Characteristics

- Separate control and data channels
  - So that control channels may be authenticated and encrypted while data channels may be as efficient as possible
- Third-party transfers (e.g., between GridFTP servers)
- Wide area network communication parameter optimization
- Integrated performance monitoring instrumentation
- Network parallel transfer streams
- Support for proxies (NAT, cache)
- Server side data striping (e.g. DPSS [7] and HPSS striped tapes)
- Server-side computation / filtering
- Support for tertiary storage system access
  - Batching of file requests
  - Tape / other near-line media pre-stage requests
  - Pinning files in the storage system cache
  - Other tape / other near-line media issues?

Issues

- Support for tertiary storage system access (as above)
  - if not in GridFTP (e.g.), then where?
  - tertiary storage cannot function correctly / optimally without this sort of functionality
- Is reliability/restart for large file transfers a minimal characteristic ? Can it be built on top of the other features? (Probably, given partial file reads.)
- We have little experience with server-side computation

- Naming abstraction is important, but can be provided by higher level services (e.g. metadata catalogues and naming services)

QoS

- Scheduling datasets to be immediately available (e.g. pre staging of tapes, pinning files in caches)

Support Required on Grid Resource Platforms

- Data access server
- Grid security infrastructure (credentials, protocol libraries)
- Grid DN based access control

Environmental Support Required

- none

Is this a minimal service?

Yes, minimal:

An essential service.

No, not minimal:

Cannot address at this time:

Current Experience

- GridFTP [8]
- SRB/MCAT [9]

3.4.2 Relational Data Base access

Is this a minimal services issue?

Yes, minimal:

- If there is a RDMS server on a Grid resource, then at least a Grid front end is probably needed for Grid authentication and access control (?)

No, not minimal:

- Can this be built on top of another Grid service? (Secure remote shell?)

Cannot address at this time:

- Do universal standard access methods currently exist?

Current Experience

- SRB/MCAT, RDBMS access [ref]
- GGF, Database Access and Integration Services Working Group
  - *Grid Data Services - Relational Database Management Systems (Version 1)*. This paper discusses issues associated with the development of relational database services, including usage scenarios. [10]
- EU DataGrid
  - *Project Spitfire - Towards Grid Web Service Databases* . This paper describes the Spitfire grid database access service for relational databases. [11]

3.4.3 Object Oriented Data Base access

Is this a minimal services issue?

Yes, minimal:

No, not minimal:

Cannot address at this time:

- No universal standard access methods currently exist (?)

Current Experience

- *Pursuit of a Scalable High Performance Multi-Petabyte Database* [12]
- *Creating Large Scale Database Servers* [13]
- *Objectivity Open File System* [14]

3.5 Asynchronous Information Sources (Events, Monitoring, Logging, etc.)

We use the term "Asynchronous Information Sources" (AIS) to mean any source of XML formatted objects that can publish its existence and object content characteristics, and then support subscription based delivery of those objects.

A model for this service is provided by the Grid Monitoring Architecture [15]. This model involves sources/producers that publish their existence and the characteristics of the data content that they supply by sending this information to a registry service. The data sinks/consumers search the registries for the desired data characteristics, and then subscribe directly with the source/producer for data delivery.

Asynchronous Information Sources include things like

- Events (system, application, workflow scripts)
- Monitoring (system parameters, batch schedulers, network, application (e.g., see [16]))
- Accounting records
- Soft state registration (of anything that changes)
- Logging of all sorts

Soft state registration, and other targeted uses (e.g. accounting, logging), should be able to be accomplished with this mechanism by, e.g.,

1. having the application implement an AIS sink/consumer function;
2. having the application implement an AIS registry function;
3. having the targeted-use source specifically register with the application sink (e.g. a logging or accounting application);
4. then have the sink automatically subscribe to everything that gets registered in its registry – i.e. all of the data sources needed for the application

Functionality

- Source registration (a la GMA, the source registers its existence and the content of the monitor objects that it will generate)
- Registry should be "globally" searchable based on various source / AIS object content characteristics
- Receiving data is by subscription and by direct transfer (source to sink) – the GMA model

Characteristics

- AIS sources should be able to register with whichever registries are appropriate, of which there may be many
- Data delivery (source to sink) should be available with multiple transport semantics
  - Streams
  - Messages
  - Unreliable multicast
  - Reliable multicast
  - Transactional
- Registry search semantics should be relational among the named AIS object fields

Issues

- Should a sink/consumer be able to "trigger" a refresh of a source/producer monitor?

QoS

- No QoS issues (?)

Support Required on Grid Resource Platforms

Basic monitoring functions with GMA source semantics need to be installed and managed.

Environmental Support Required

Registry servers must be provided and maintained.

Is this a minimal service?

Yes, minimal:

- Cannot, in general, be built using the job initiation service
  - + Users cannot, in general, start and maintain long lived servers on Grid resources
  - + There will be Grid system that we need to monitor, but which will not support a job initiation service (e.g. instruments and storage systems)

No, not minimal:

Cannot address at this time:

Current Experience

- GGF GMA docs
- Implementations
  - *Distributed Monitoring Framework (DMF),* [17]
  - *Information and Monitoring Services Architecture*, [18]
  - *A Framework for Control and Observation in Distributed Environments*, [19]
- Examples
  - *Monitoring Data Archives for Grid Environments*, [16]

3.6 Remote Authentication, Delegation, and Secure Communication

Remote authentication is accomplished by techniques that verify a cryptographic identity in a way that establishes trust[1] in an unbroken chain from the relying party back to a named human, system, or service identity. This is accomplished in a sequence of trusted steps, each one of which is essential in order to get from accepting a remote user on a Grid resource back to a named entity.

Delegation involves generating and sending a proxy certificate and its private key to a remote Grid system so that remote system may act on behalf of the user. This is the essence of the single sing-on provided by the Grid: A user / entity proves its identity once, and then delegates its authority to remote systems for subsequent processing steps.

A secure communication channel is derived from the Public Key Infrastructure process and the IETF Transport Level Security protocol, as described below.

The trust establishment process involves:

1. Binding an entity identity to a Distinguished Name ("DN" - the subject name in an X.509 identity certificate)
   - Trust in this step is accomplished through the (published and audited) policy based identity verification procedures of the Certification Authority that issues the identity certificates
2. Binding a public key to the DN (generating an X.509 certificate)
   - Trust in this step is accomplished through the (published and audited) policy based operational procedures of the issuing Certification Authority ("CA").
3. Assurance that the public key that is presented actually represents the user
   - Trust in this step comes from the cryptography and protocols of Public Key Infrastructure.
4. Assurance that a message tied to the entity DN could only have originated with that entity:
   - Trust that a message signed by a private key could only have been signed by the private key corresponding to the public key (and therefore the named entity via X.509 certs) comes from public key cryptography
   - Trust in this step is also through user key management (the mechanism by which the user limits the use of its identity), which is assured by user education, care in dealing with one's cyber environment, and shared understanding as to the significance of the private key.
5. Mutual authentication, whereby two ends of a communication channel agree on each other's identity
   - Trust in this step is through the cryptographic techniques and protocols of the Transport Level Security ("TLS") standard.
6. Delegation of identity to remote Grid systems
   - Trust in this step is through the cryptographic techniques and protocols for generating, managing, and using proxy certificates that are directly derived from the CA issued identity certificates.

At this point a cryptographic proxy identity is present at the remote Grid resources, and it is trusted to represent the original named entity.

By virtue of using TLS for mutual authentication, a secure communication channel has been established at this point. The channel uses symmetric key cryptography, and session key management for this secure channel is part of the TLS protocol.

Confidentiality is a characteristic of the TLS channel.

[1] Trust is "confidence in or reliance on some quality or attribute of a person or thing, or the truth of a statement." Oxford English Dictionary, Second Edition (1989). Oxford University Press.

Message integrity comes with the TLS channel, but not without encryption. The GSS-API provides an integrity-only function.

3.6.1 Certification Authority and Certificate Management

The CA accomplishes steps 1 and 2 in the trust establishment process.

CAs have clear and published policy on the circumstances under which they will issue identity certificates, and how the DNs are generated. While this policy is not uniform, it is public so that each virtual or actual organization that supplies resources to the Grid may make an informed decision on whether to accept the remote user identity certificates or not.

CAs have clear and published policy on there operating procedures that indicate the level of care taken in the certificate generating and issuing process in order to ensure that certificates traceable back to the CA are not forged or otherwise cryptographically compromised.

Functionality

- Provides a mechanism for users / entities to request certificates
- Provides a registration process that verifies user/entity identity
- Issues and signs X.509 identity certificates
- Provides Certificate Revocation List generation, management, access, and use
- Provides a certificate repository

Characteristics

- Operated according to a formal, published policy statement, and with some level of audit
- Highly secure
- Formal logging

Issues

- Negotiating common policy among multi-institutional and/or international VOs is hard
- Is a certificate repository a minimal function?
- Certificate Revocation List distribution protocols and management are not well defined yet

QoS

- No QoS issues (?)

Support Required on Grid Resource Platforms

- The CA public key must be installed on all platforms that authenticate Grid entities via the Grid security infrastructure – for a mixed user population this will involve multiple CAs
- The CA signing policy file must be installed on all platforms that authenticate Grid entities via the Grid security infrastructure (this allows a relying party to restrict the certs that it will accept from a given CA based on the name space of the DN)

Environmental Support Required

- A secure physical and cyber infrastructure are needed for the CA

Is this a minimal service?

Yes, minimal:

- The Certification Authority, and its published policies, are an essential component for this service for establishing trust in remote user access

No, not minimal:

Cannot address at this time:

Current Experience

- See the DOE Science Grid CA (doegrids.org, [20]) for an example of a production, multiple VO CA
  - See the project site for the DOE Science Grid CA (http://envisage.es.net) which has the project documents and history [21]
- EU DataGrid
  - Certification Authorities – see [22]
  - Cross certification check list – see [23]

3.6.2 User Key Management

The user / entity private key corresponding to the public key that the CA has bound to the DN is what the remote entity uses to prove that it is the entity represented in the certificate. This key represents the user identity and must be carefully protected. The CA must convey the importance of protecting the private key to the user when the certificate is issued.

At least one cryptographic mechanism must be provided for protecting the private key. Typical mechanisms are to keep the private key encrypted with a second (symmetric) key that is protected either with a passphrase, or keep it in a cryptographic device such as a smart card.

3.6.3 Mutual Authentication

The Transport Level Security protocol (formerly know as SSL) uses a cryptographic protocol, together the tracability of the certificates back to an issuing authority that is policy based, to establish an authenticated and secure communication channel.

Functionality

- Provides for using host identity credentials at both ends of a transport connection for
  - validating the system identities
  - conveying user / Grid entity credentials to the remote system

Characteristics

- Provided by IETF TLS

Issues

QoS

- No QoS issues (?)

Support Required on Grid Resource Platforms

- Host identity certificates must be installed on the Grid resource systems
- Host certificates must be issued by a CA according to a clear policy
- Compatible cryptographic libraries must be installed on the receivers (e.g. OpenSSL)

Environmental Support Required

- Care must be taken to keep the TLS implementation up-to-date. Several vulnerabilities have been detected and corrected in the recent past.

Is this a minimal service?

Yes, minimal:

This is an essential component for the Grid security service to securely access a remote system

No, not minimal:

Cannot address at this time:

Current Experience

- Globus GSI [24]
- OpenSSL [25]

3.6.4 Secure Communication

Following mutual authentication, the Transport Level Security protocol (formerly know as SSL) uses a cryptographic protocol to establish a secure communication channel.

Functionality

- A secure, stream oriented communication

Characteristics

- Provided by IETF TLS

Issues

- As above (section 3.6.3, “Mutual Authentication”)

QoS

- As above (section 3.6.3, “Mutual Authentication”)

Support Required on Grid Resource Platforms

- As above (section 3.6.3, “Mutual Authentication”)

Environmental Support Required

- As above (section 3.6.3, “Mutual Authentication”)

Is this a minimal service?

Yes, minimal:

As above (section 3.6.3, “Mutual Authentication”)

No, not minimal:

Cannot address at this time:

Current Experience

- As above (section 3.6.3, “Mutual Authentication”)

3.6.5 Delegation

Delegation is the process by which the user's identity is carried to a remote system without the user being directly involved at the remote system.

This involves generating a proxy certificate that is derived from the user's identity certificate. The proxy, its private key, and the user identity certificate are all conveyed to the remote systems in order to support authentication and authorization, and to support conveying the user identity to subsequent systems that may be needed for the Grid task.

This is accomplished with a delegation protocol.

Functionality

- Generates "proxy" certificates
- Provides a protocol for conveying the proxy to the remote site

Characteristics

Issues

- Clearly defined API for delegation
- Limited delegation

QoS

- No QoS issues (?)

Support Required on Grid Resource Platforms

- Software support for the delegation process

Environmental Support Required

- The private key of the proxy is stored in a file that should be "well protected." E.g., it should only be readable by the UID of the process that must use that key to generate downstream proxies.

Is this a minimal service?

Yes, minimal:

Delegation is an essential Grid service. For a complex Grid task operating in a large Grid, it would be virtually impossible for the user to directly interact with every system that might be involved.

No, not minimal:

Cannot address at this time:

Current Experience

- Globus GSI [24]
- Draft TLS extensions for delegation [26]
- GSS-API extensions for Grids [27]
- GSI Online Credential Retrieval – Requirements [28]

3.6.6 GSS-API

The IETF GSS-API provides an API for security context establishment, message integrity, and message confidentiality.

Functionality

- An API for security context establishment, message integrity, and message confidentiality

Characteristics

Issues

- GSS-API implementations also do framing, which makes GSS-API more then just an API. This means that the receiver must understand the GSS framing.

QoS

- No QoS issues (?)

Support Required on Grid Resource Platforms

- GSS libraries

Environmental Support Required

- none

Is this a minimal service?

Yes, minimal:

- No other GSI service provides message integrity without encryption

No, not minimal:

- Is message integrity w/o encryption essential?

Cannot address at this time:

Current Experience

- Globus GSI [24]
- GSS API [29]

3.6.7 The Overall Grid Security Infrastructure Service

Functionality

- As above

Characteristics

Issues

- Well defined APIs for the basic functions
- GSS-API issue as above
- Interfaces to local security domains that are not PKI, e.g. Kerberos
- On-line credential repositories
- CRLs
- Is basic ACL (access control list – e.g. the Globus mapfile) authorization a minimal function?

QoS

- No QoS issues (?)

Support Required on Grid Resource Platforms

- See individual components
- ACLs for authorization

Environmental Support Required

- See individual components

Is this a minimal service?

Yes, minimal:

- Secure authentication and communication are essential Grid functions
- Trust by the relying party (the remote Grid resource) in the identity of the entity seeking to establish a secure communication channel is an essential service. This service actually has several different components, all of which are required to provide the service.
- Secure, authenticated transport is essential for Grid service command channels / messages

No, not minimal:

- Authorization is a site policy issue

Cannot address at this time:

Current Experience

- PKI [30]
- Globus implementation of the Grid Security Infrastructure [24]
- Globus implementation of GSS-API [24]

3.7 System Management and Access

System management, and sometimes remote user access, are needed so that Grid resources may be managed and interactively accessed within the Grid context.

Functionality

- Remote login, authenticated and secured with Grid security functions and authorization based on a Grid DN
- Remote shell, authenticated and secured with Grid security functions and authorization based on a Grid DN
- Remote copy, authenticated and secured with Grid security functions and authorization based on a Grid DN

Characteristics

Issues

-

QoS

- No QoS issues (?)

Support Required on Grid Resource Platforms

- The servers that support these services must be installed and maintained.
- Host certificates must be installed and managed
- User CA keys must be installed and managed
- ACLs for authorization must be installed and managed

Environmental Support Required

- none

Is this a minimal service?

Yes, minimal:

- This seems to be an essential service, because if it is not provided then this is always accomplished in a ad-hoc manner.

No, not minimal:

Cannot address at this time:

Current Experience

- GSIssh [31]

3.8 Architectural Constraints

In order to be called a Grid Common Service, it should not be possible to convey command and control messages to remote Grid systems except through the secure and authenticated communication provided by the Grid security functions. This is indicated pictorially in Figure 1, “Core Grid Functions Relationships.”

A Grid without this sort of security is not a Grid.

Secure data channels should always be optional, as the may be impractical in some circumstances. Should data channels always be authenticated, but not encrypted? (Is this possible with TLS?)

Others?

3.9 Bindings

Most of the Core Functions will be defined in terms of protocols and data structures, and this provides the basic uniformity required of Grids.

However, there will be many ways to use these Core Functions. For example

- Globus toolkit's C language [, 2001 #72]
- CoG kit's Java interface to the Globus functions [Laszewski, 2001 #69]
- PyGlobus interface to the Globus functions [Jackson, 2002 #194]
- Arguably the OGSI work [Global Grid Forum, 2002 #191] represents a non-Globus interface to the Core functions

And there will be others.

Issues

It may be desirable, even necessary, to require that a minimal implementation of the core functions include the client side bindings for a representative set of programming styles. For example, the once given above: C, Java, Python, OGSI.

3.10 Other / Future Services as Core Grid Functions

These functions are noted for future discussion, but are not current practice in Grids.

Such functions might be generated by

- Classes of Grid resources that have not yet been considered
- New application types and/or uses that require new core functionality

Recall that criteria for core / minimal Grid functions are:

1. it be an important function
2. it cannot be built from existing Grid services, and therefore requires some operational elements on Grid resources (servers, libraries, etc.)

3.10.1 Abstraction of Computing Resource Architecture

Provide for a mechanism to map a workflow onto different computing resource architectures.

Account, e.g., for how data is staged into and out of systems, how directory structures are set up for multiple tasks, how data is cached from one task to the next, etc.

UNICORE [32] addresses this issue.

3.10.2 Transactional Messaging

Is transactional messaging a core Grid function?

3.10.3 Reliable, Secure Multicast

Is reliable, secure multicast / secure group communication (see, e.g., [33]) a core Grid function?

3.10.4 Checkpoint / Restart / Coordinated Recovery

Is checkpoint / restart / coordinated recovery a core Grid function?

3.10.5 Structured Data Access

Are access methods for

- XML objects
- Time series

core Grid functions?

3.10.6 Quality of Service

- Is this part of the service request?
- Part of scheduling?
- Separate?
- More than scheduling?

3.10.7 Debug

Any special support needed for Grid debugging?

- Seems like local helper processes that are initiated by Grid job initiator will serve this purpose.
- Is there a permissions issue (how do I debug someone else's job?)
- Need to talk with Bob Hood (rhood@nas.nasa.gov) about this (he has built a Grid debugger)

3.10.8 Communications channel "tapping"

- for debug, steering, analysis
- where is Nexus when we need it?

3.10.9 Authorization

Issue: Is some type of access control a requirement? Is this a local control issue?

- Access control lists
- CAS [34]
- Attribute certificate based [35]

4. Security Considerations

Security is a fundamental aspect of this document, which describes the relationship of security to the core Grid functions.

5. Glossary

6. Author Contact Information

William E. Johnston
Lawrence Berkeley National Laboratory / NASA Ames Research Center
tel: +1-510-486-5014, fax: +1-603-719-1356
USMail: 1 Cyclotron Rd., MS 50B-2239, Berkeley, CA, 94720, USA
wejohnston@lbl.gov

John M. Brooke
Manchester Research Centre for Computational Science (MRCCS)
Manchester Computing, University of Manchester
Manchester M13 9PL, UK
Tel: +44 (0)161 275 6814 Fax: +44 (0)161 275 6800 Mobile 07765 220 227
http://www.csar.cfs.ac.uk/staff/brooke Email: j.m.brooke@man.ac.uk

7. Acknowledgements

8. Notices

8.1 Intellectual Property Statement

The GGF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the GGF Secretariat.

The GGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this recommendation. Please address the information to the GGF Executive Director (see contacts information at GGF website).

8.2 Full Copyright Notice

9. Notes and References

[1] **Grid Information Services / MDS**, Globus Project. http://www.globus.org/mds/
Grid computing technologies enable wide-spread sharing and coordinated use of networked resources. Sharing relationships may be static and long-lived—e.g., among the major resource centers of a company or university—or highly dynamic: e.g., among the evolving membership of a scientific collaboration. In either case, the fact that users typically have little or no knowledge of the resources contributed by participants in the “virtual organization” (VO) poses a significant obstacle to their use. For this reason, information services designed to support the initial discovery and ongoing monitoring of the existence and characteristics of resources, services, computations, and other entities are a vital part of a Grid system. ("Grid Information Services for Distributed Resource Sharing" - http://www.globus.org/research/papers/MDS-HPDC.pdf)
The Monitoring and Discovery Service architecture addresses the unique requirements of Grid environments. Its architecture consists of two basic elements:
- A large, distributed collection of generic information providers provide access to information about individual entities, via local operations or gateways to other information sources (e.g., SNMP queries). Information is structured in term of a standard data model, taken from LDAP: an entity is described by a set of "objects" comprised of typed attribute-value pairs.
- Higher-level services, collect, manage, index, and/or respond to information provided by one or more information providers. We distinguish in particular aggregate directory services, which facilitate resource discovery and monitoring for VOs by implementing both generic and specialized views and search methods for a collection of resources. Other higher-level services can use this information and/or information obtained directly from providers for the purposes of brokering, monitoring, troubleshooting, etc.
Interactions between higher-level services (or users) and providers are defined in terms of two basic protocols: a soft-state registration protocol for identifying entities participating in the information service, and an enquiry protocol for retrieval of information about those entities, whether via query or subscription. In brief, a provider uses the registration protocol to notify higher-level services of its existence; a higher-level service uses the enquiry protocol to obtain

information about the entities known to a provider, which it merges into its aggregate view. Integration with the Grid Security Infrastructure (GSI) provides for authentication and access control to information.

[2] **Super Scheduler Steps/Framework**, J. Schopf. http://www.mcs.anl.gov/~jms/ggf-sched/WD/schedwd.8.5.doc
http://www.mcs.anl.gov/~jms/ggf-sched/WD/schedwd.8.5.pdf
Overview of current user practices for scheduling across administrative domains. GGF document.

[3] **SNAP: A Protocol for Negotiating Service Level Agreements and Coordinating Resource Management in Distributed Systems**, K. Czajkowski, I. Foster, V. Sander, C. Kesselman and S. Tuecke. In *8th Workshop on Job Scheduling Strategies for Parallel Processing*. 2002. Edinburgh, Scotland.http://www.globus.org/research/papers/jsspp02-snap-preprint.pdf
A fundamental problem with distributed applications is to map activities such as computation or data transfer onto a set of resources that will meet the application's requirement for performance, cost, security, or other quality of service metrics. An application or client must engage in a multi-phase negotiation process with resource managers, as it discovers, reserves, acquires, configures, monitors, and potentially renegotiates resource access. Current approaches to resource management tend to specialize for specific classes of resource (processor, network, etc.), and have addressed coordination across resources in a limited fashion, if at all. We present a generalized resource management model in which resource interactions are mapped onto a well defined set of platform-independent service level agreements (SLAs). We instantiate this model in the Service Negotiation and Acquisition Protocol (SNAP) which provides lifetime management and an at-most-once creation semantics for remote SLAs. The result is a resource management framework for distributed systems that we believe is more powerful and general than current approaches. We explain how SNAP can be deployed within the context of the Globus Toolkit.

[4] **The Portable Batch Scheduler**. http://www.pbspro.com/tech_overview.html
The purpose of the PBS system is to provide additional controls over initiating or scheduling execution of batch jobs; and to allow routing of those jobs between different hosts [that run administratively coupled instances of PBS]. The batch system allows a site to define and implement policy as to what types of resources and how much of each resource can be used by different jobs. The batch system also provides a mechanism with which a user can insure a job will have access to the resources required to complete.

[5] **Maui Silver Metascheduler**.
http://www.supercluster.org/documentation/silver/silveroverview.html
Silver is an advance reservation metascheduler. Its design allows it to load balance workload across multiple systems in completely independent administrative domains. How much or how little a system participates in this load sharing activity is completely up to the local administration. All workload is tracked and accounted for allowing 'allocation' exchanges to take place between the active sites.

[6] **Globus Resource Allocation Manager (GRAM)**, Globus Project. 2002. http://www-fp.globus.org/gram/overview.html
The Globus Resource Allocation Manager (GRAM) is the lowest level of Globus resource management architecture. GRAM allows you to run jobs remotely, providing an API for submitting, monitoring, and terminating your job.
To run a job remotely, a GRAM gatekeeper (server) must be running on a remote computer, listening at a port; and the application needs to be compiled on that remote machine. The execution begins when a GRAM user application runs on the local machine, sending a job request to the remote computer.
The request is sent to the gatekeeper of the remote computer. The gatekeeper handles the request and creates a job manager for the job. The job manager starts and monitors the remote program, communicating state changes back to the user on the local machine. When the remote application terminates, normally or by failing, the job manager terminates as well.
The executable, stdin and stdout, as well as the name and port of the remote computer, are specified as part of the job request. The job request is handled by the gatekeeper, which creates

a job manager for the new job. The job manager handles the execution of the job, as well as any communication with the user.

[7] **A Network-Aware Distributed Storage Cache for Data Intensive Environments**, B. Tierney, J. Lee, B. Crowley, M. Holding, J. Hylton and F. Drake. In *Proc. 8th IEEE Symp. on High Performance Distributed Computing*. 1999.http://www-didc.lbl.gov/papers/dpss.hpdc99.pdf

[8] **The GridFTP Protocol and Software**, Globus Project. 2002. http://www.globus.org/datagrid/gridftp.html
GridFTP is a high-performance, secure, reliable data transfer protocol optimized for high-bandwidth wide-area networks. The GridFTP protocol is based on FTP, the highly-popular Internet file transfer protocol. We have selected a set of protocol features and extensions defined already in IETF RFCs and added a few additional features to meet requirement from current data grid projects.

[9] **The Storage Resource Broker**. http://www.npaci.edu/DICE/SRB/
The SDSC Storage Resource Broker (SRB) is a client-server middleware that provides a uniform interface for connecting to heterogeneous data resources over a network and accessing replicated data sets. SRB, in conjunction with the Metadata Catalog (MCAT), provides a way to access data sets and resources based on their attributes rather than their names or physical locations.

[10] **Grid Data Services - Relational Database Management**, B. Collins, A. Borley, N. Hardman, A. Knox, S. Laws, J. Magowan, M. Oevers and E. Zaluska. http://www.cs.man.ac.uk/grid-db/papers/grdb.pdf
This paper discusses issues associated with the development of relational database services, including usage scenarios.

[11] **Project Spitfite - Towards Grid Web Service Databases**, W. H. Bell, D. Bosio, W. Hoschek, P. Kunszt, G. McCance and M. Silander. 2002. http://www.cs.man.ac.uk/grid-db/papers/ggf5-spitfire.pdf
This paper describes the Spitfire grid database access service for relational databases. This is an EU DataGrid project. See http://hep-proj-spitfire.web.cern.ch/hep-proj-spitfire/server/doc/

[12] **Pursuit of a Scalable High Performance Multi-Petabyte Database**, A. Hanushevsky and M. Nowak. In *Sixteenth IEEE Mass Storage Systems Symposium*. 1999.http://www.slac.stanford.edu/BFROOT/www/Public/Computing/Databases/proceedings/ieee_16mssc99.pdf
When the BaBar experiment at the Stanford Linear Accelerator Center starts in April 1999, it will generate approximately 200TB/year of data at a rate of 10MB/sec for 10 years. A mere six years later, CERN, the European Laboratory for Particle Physics, will start an experiment whose data storage requirements are two orders of magnitude larger. In both experiments, all of the data will reside in Objectivity databases accessible via the Advanced Multi-threaded Server (AMS). The quantity and rate at which the data is produced requires the use of a high performance hierarchical mass storage system in place of a standard Unix file system. Furthermore, the distributed nature of the experiment, involving scientists from 80 Institutions in 10 countries, also requires an extended security infrastructure not commonly found in standard Unix file systems. The combination of challenges that must be overcome in order to effectively deal with a multi-petabyte object oriented database is substantial. Our particular approach marries an optimized Unix file system with an industrial strength Mass Storage System. This paper describes what we had to do to create a robust and uniform system based on these components.

[13] **Creating Large Scale Database Servers**, J. Becla and A. Hanushevsky. In *Ninth IEEE International Symposium on High Performance Distributed Computing (HPDC'00)*. 2000.http://www.slac.stanford.edu/BFROOT/www/Public/Computing/Databases/proceedings/hpdc2000.pdf
The BaBar experiment at the Stanford Linear Accelerator Center (SLAC) is designed to perform a high precision investigation of the decays of the B-meson produced from electron-positron interactions. The experiment, started in May 1999, will generate approximately 300TB/year of data for 10 years. All of the data will reside in Objectivity databases accessible via the Advanced Multi-threaded Server (AMS). To date, over 70TB of data have been placed in Objectivity/DB,

making it one of the largest databases in the world. Providing access to such a large quantity of data through a database server is a daunting task. A full-scale testbed environment had to be developed to tune various software parameters and a fundamental change had to occur in the AMS architecture to allow it to scale past several hundred terabytes of data. Additionally, several protocol extensions had to be implemented to provide practical access to large quantities of data. This paper will describe the design of the database, the changes that we needed to make in the AMS for scalability reasons, and how the lessons we learned would be applicable to virtually any kind of database server seeking to operate in the Petabyte region.

[14] **Objectivity Open File System**, A. Hanushevsky. In *HEPNT-HEPiX fall '99*. 1999.http://www-project.slac.stanford.edu/hepix/HEPiX99-oofs.ppt
Describes the architecture of a system that integrates Objectivity's OODB with HPSS.

[15] **Grid Monitoring Architecture Working Group**, Global Grid Forum. http://www-didc.lbl.gov/GGF-PERF/GMA-WG/
The Grid Monitoring Architecture working group is focused on producing a high-level architecture statement of the components and interfaces needed to promote interoperability between heterogeneous monitoring systems on the Grid. The main products of this work are the architecture document itself, and accompanying case studies that illustrate the concrete application of the architecture to monitoring problems.

[16] **Monitoring Data Archives for Grid Environments**, J. Lee, D. Gunter, M. Stoufer and B. Tierney. In *SC2002*. 2002.http://www-didc.lbl.gov/publications.html
Developers and users of high-performance distributed systems often observe performance problems such as unexpectedly low throughput or high latency. To determine the source of these performance problems, detailed end-to-end monitoring data from applications, networks, operating systems, and hardware must be correlated across time and space. Researchers need to be able to view and compare this very detailed monitoring data from a variety of angles. To solve this problem, we propose a relational monitoring data archive that is designed to efficiently handle high-volume streams of monitoring data. In this paper we present an instrumentation and event archive service that can be used to collect and aggregate detailed end-to-end monitoring information from distributed applications. This archive service is designed to be scalable and fault tolerant. We also show how the archive is based on "Grid Monitoring Architecture" defined by the Global Grid Forum.

[17] **Distributed Monitoring Framework (DMF)**, Lawrence Berkeley National Lab. http://www-didc.lbl.gov/DMF/
The goal of the Distributed Monitoring Framework is to improve end-to-end data throughput for data intensive applications in a high-speed WAN environments, and to provide the ability to do performance analysis and fault detection in a Grid computing environment. This monitoring framework will provide accurate, detailed, and adaptive monitoring of all of distributed computing components, including the network. Analysis tools will be able to use this monitoring data for real-time analysis, anomaly identification, and response.
Many of the components of the DMF have already been prototyped or implemented by the DIDC Group. The NetLogger Toolkit includes application sensors, some system and network sensors, a powerful event visualization tool, and a simple event archive. The Network characterization Service has proven to be a very useful hop-by-hop network sensor. Our work on the Global Grid Forum Grid Monitoring Architecture (GMA) addressed the event management system. JAMM (Java Agents for Monitoring Management) is preliminary work on sensor management. The Enable project produced a simple network tuning advice service.

[18] **Information and Monitoring Services Architecture**, European Union DataGrid - WP3. http://hepunx.rl.ac.uk/edg/wp3/documentation/doc/arch/index.html
The aim of this work package is to specify, develop, integrate and test tools and infrastructure to enable end-user and administrator access to status and error information in a Grid environment and to provide an environment in which application monitoring can be carried out. This will permit both job performance optimisation as well as allowing for problem tracing and is crucial to facilitating high performance Grid computing.

[19] **A Framework for Control and Observation in Distributed Environments**, W. Smith. NASA Ames Research Center. http://www.nas.nasa.gov/~wwsmith/papers.html
A GGF, GMA implementation.

[20] **DOE Science Grid PKI Certificate Policy And Certification Practice Statement**. http://www.doegrids.org/
This document represents the policy for the DOE Science Grid Certification Authority operated by ESnet. It addresses Certificate Policy (CP) and Certification Practice Statement (CPS). The CP is a named set of rules that indicates the applicability of a certificate to a particular community and/or class of application with common security requirements. For example, a particular certificate policy might indicate applicability of a type of certificate to the authentication of electronic data interchange transactions for the trading of goods within a given price range. The CPS is a statement of the practices, which a certification authority employs in issuing certificates.
[21] **ESnet's SciDAC PKI & Directory Project - Homepage**, T. Genovese and M. Helm. DOE Energy Sciences Network. http://envisage.es.net/
This is the ESnet PKI project site. ESnet is building a Public Key Infrastructure service to support the DOE Science Grid, SciDAC projects and other DOE research efforts. The main goal is to provide DOE scientist and engineers Identity and Service certificates that allow them to participate in the growing national and international computational Grids.
[22] **Certification Authorities**, European Union DataGrid. 2002. http://marianne.in2p3.fr/datagrid/ca/ca-table-ca.html
The current list of EU DataGrid recognized CAs and their certificates.
[23] **Certification Authorities Acceptance and Feature Matrices**, European Union DataGrid. 2002. http://www.cs.tcd.ie/coghlan/cps-matrix/
The Acceptance and Feature matrices are key aspects of establishing cross-site trust.
[24] **Grid Security Infrastructure (GSI)**, Globus Project. http://www.globus.org/security/
The primary elements of the GSI are identity certificates, mutual authentication, confidential communication, delegation, and single sign-on.
GSI is based on public key encryption, X.509 certificates, and the Secure Sockets Layer (SSL) communication protocol. Extensions to these standards have been added for single sign-on and delegation. The Globus Toolkit's implementation of the GSI adheres to the Generic Security Service API (GSS-API), which is a standard API for security systems promoted by the Internet Engineering Task Force (IETF).
[25] **OpenSSL**. http://www.openssl.org/
The OpenSSL Project is a collaborative effort to develop a robust, commercial-grade, full-featured, and Open Source toolkit implementing the Secure Sockets Layer (SSL v2/v3) and Transport Layer Security (TLS v1) protocols as well as a full-strength general purpose cryptography library. The project is managed by a worldwide community of volunteers that use the Internet to communicate, plan, and develop the OpenSSL toolkit and its related documentation.
[26] **Internet X.509 Public Key Infrastructure Proxy Certificate Profile**, S. Tuecke, D. Engert, I. Foster, V. Welch, M. Thompson, L. Pearlman and C. Kesselman. February 2002. http://www.gridforum.org/security/ggf4_2002-02/draft-ietf-pkix-proxy-02.txt http://www.gridforum.org/security/ggf4_2002-02/draft-ietf-pkix-proxy-02.pdf
A technical specification draft of the X.509 certificate extensions required to support proxies, which is used for GSI single sign-on and delegation.
[27] **GSS-API Extensions**, S. Meder, V. Welch, S. Tuecke and D. Engert. February 2002. http://www.gridforum.org/security/ggf4_2002-02/draft-ggf-gss-extensions-05.doc http://www.gridforum.org/security/ggf4_2002-02/draft-ggf-gss-extensions-05.pdf
This document defines extensions to RFC 2743, Generic Security Service Application Program Interface Version 2, Update 1. Extensions include: credential export and import of credentials; delegation at any time; credential extensions (e.g. restrictions) handling.
[28] **GSI Online Credential Retrieval - Requirements**, J. Basney. February 2002. http://www.gridforum.org/security/ggf4_2002-02/draft-ggf-gsi-ocr-requirements-01.doc http://www.gridforum.org/security/ggf4_2002-02/draft-ggf-gsi-ocr-requirements-01.pdf
An online credential retrieval (OCR) service gives users secure and convenient access to the credentials they need for authentication. To make credentials available, the service either stores the credentials in a secure repository or generates new credentials on request.
This memo defines requirements for online credential retrieval services that provide secure access to X.509 credentials in the Grid Security Infrastructure (GSI).

[29] **Generic Security Service Application Program Interface, Version 2**, J. Linn. http://www.ietf.org/rfc/rfc2078.txt?number=2078
The Generic Security Service Application Program Interface (GSS-API), as defined in RFC-1508, provides security services to callers in a generic fashion, supportable with a range of underlying mechanisms and technologies and hence allowing source-level portability of applications to different environments. This specification defines GSS-API services and primitives at a level independent of underlying mechanism and programming language environment, and is to be complemented by other, related specifications:
documents defining specific parameter bindings for particular language environments
documents defining token formats, protocols, and procedures to be implemented in order to realize GSS-API services atop particular security mechanisms
[30] PKI.
Public-Key certificate infrastructure ("PKI") provides the tools to create and manage digitally signed certificates. For identity authentication, a certification authority generates a certificate (most commonly an X.509 certificate) containing the name (usually X.500 distinguished name) of an entity (e.g. user) and that entity's public key. The CA then signs this "certificate" and publishes it (usually in an LDAP directory service). These are the basic components of PKI, and allow the entity to prove its identity, independent of location or system. For more information, see, e.g., RSA Lab's "Frequently Asked Questions About Today's Cryptography" http://www.rsa.com/rsalabs/faq/, Computer Communications Security: Principles, Standards, Protocols, and Techniques. W. Ford, Prentice-Hall, Englewood Cliffs, New Jersey, 07632, 1995, or Applied Cryptography, B. Schneier, John Wiley & Sons, 1996.
[31] **GSI-Enabled OpenSSH**, NCSA. http://www.ncsa.uiuc.edu/Divisions/ACES/GSI/openssh/
NCSA maintains a patch to OpenSSH that adds support for GSI authentication.
[32] **UNICORE**. http://www.unicore.de/
UNICORE lets the user prepare or modify structured jobs through a graphical user interface on a local Unix workstation or a Windows PC. Jobs can be submitted to any of the platforms of a UNICORE GRID and the user can monitor and control the submitted jobs through the job monitor part of the client.
A UNICORE job contains a number of interdependent tasks. The dependencies indicate temporal relations or data transfer. Currently, execution of scripts, compile, link, execute tasks and data transfer directives are supported. An execution system request associated with a job specifies where its tasks are to be run. Tasks can be grouped into sub-jobs, creating a hierarchical job structure and allowing different steps to execute on different systems within the UNICORE GRID.
[33] **Reliable and Secure Group Communication**, D. Agarwal, K. Berket and O. Chevassut. http://www-itg.lbl.gov/CIF/GroupComm

[34] **Community Authorization Service (CAS)**, Globus Project. 2002. http://www.globus.org/security/CAS/
CAS allows resource providers to specify course-grained access control policies in terms of communities as a whole, delegating fine-grained access control policy management to the community itself. Resource providers maintain ultimate authority over their resources but are spared day-to-day policy administration tasks (e.g. adding and deleting users, modifying user privileges). Briefly, the process is: 1) A CAS server is initiated for a community: a community representative acquires a GSI credential to represent that community as a whole, and then runs a CAS server using that community identity. 2) Resource providers grant privileges to the community. Each resource provider verifies that the holder of the community credential represents that community and that the community's policies are compatible with the resource provider's own policies. Once a trust relationship has been established, the resource provider then grants rights to the community identity, using normal local mechanisms (e.g. gridmap files and disk quotas, filesystem permissions, etc.). 3) Community representatives use the CAS to manage the community's trust relationships (e.g., to enroll users and resource providers into the community according to the community's standards) and grant fine-grained access control to resources. The CAS server is also used to manage its own access control policies; for example, community members who have the appropriate privileges may authorize additional community members to manage groups, grant permissions on some or all of the community's resources, etc.

4) When a user wants to access resources served by the CAS, that user makes a request to the CAS server. If the CAS server's database indicates that the user has the appropriate privileges, the CAS issues the user a GSI restricted proxy credential with an embedded policy giving the user the right to perform the requested actions. 5) The user then uses the credentials from the CAS to connect to the resource with any normal Globus tool (e.g. GridFTP). The resource then applies its local policy to determine the amount of access granted to the community, and further restricts that access based on the policy in the CAS credentials, This serves to limit the user's privileges to the intersection of those granted by the CAS to the user and those granted by the resource provider to the community.

[35] **Certificate-based Access Control for Widely Distributed Resources**, M. Thompson, W. Johnston, S. Mudumbai, G. Hoo, K. Jackson and A. Essiari. In *Eighth Usenix Security Symposium*. 1999.http://www-itg.lbl.gov/Akenti/papers.html